

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

Claim 1. (original): A method for multi-site clustering in a network, comprising:
receiving a domain name from a client at a first domain name server;
identifying a first network address and a second network address associated with the domain name, the first network address associated with a first site that includes a load balancer coupled to a plurality of web servers, the second network address associated with a second site that includes a second domain name server;
determining whether the first site is available to serve the client;
determining whether the second site is available to serve the client;
communicating the first network address followed by the second network address to the client based at least partially on a determination that the first site and the second site are available; and
communicating the second network address to the client based at least partially on a determination that the first site is not available.

Claim 2. (original): The method of Claim 1, wherein determining whether the first site is available comprises determining whether at least one of the web servers is available.

Claim 3. (original): The method of Claim 2, wherein determining whether at least one of the web servers is available comprises:
communicating a message to the web servers; and
determining whether the web servers respond to the message.

Claim 4. (original): The method of Claim 1, wherein:
the web servers are coupled to a database; and
determining whether the first site is available comprises determining whether the database is available.

Claim 5. (original): The method of Claim 4, wherein determining whether the database is available comprises:

instructing at least one of the web servers to execute a web page;
receiving the executed web page; and
determining if the executed web page contains at least one expected keyword.

Claim 6. (original): The method of Claim 5, wherein the expected keyword comprises a name of a database server serving the database.

Claim 7. (original): The method of Claim 1, wherein determining whether the second site is available comprises determining whether the second domain name server is available.

Claim 8. (original): The method of Claim 7, wherein determining whether the second domain name server is available comprises determining if at least one message has been received at the first domain name server from the second domain name server during each of a plurality of time periods.

Claim 9. (original): The method of Claim 1, further comprising communicating the first network address to the client when the first site is available and the second site is not available.

Claim 10. (original): The method of Claim 1, wherein communicating the first network address followed by the second network address to the client is also based on at least one of a load placed on each of the first and second sites, a response time of each of the first and second sites, and a locality of the client to each of the first and second sites.

Claim 11. (original): The method of Claim 10, further comprising communicating the second network address followed by the first network address to the client based on the determination that the first site and the second site are available and at least one of the load placed on each of the first and second sites, the response time of each of the first and second sites, and the locality of the client to each of the first and second sites.

Claim 12. (original): The method of Claim 1, wherein: the first site includes a first database; and the second site includes a second database.

Claim 13. (original): The method of Claim 12, further comprising synchronizing information in the first and second databases.

Claim 14. (original): The method of Claim 13, wherein synchronizing the information in the first and second databases comprises:

- making first changes to the information in the first database;
- communicating the first changes to the second database;
- receiving second changes from the second database; and
- replicating the second changes to the information in the first database.

Claim 15. (original): The method of Claim 14, wherein each first change comprises one of an addition of information to the first database, a modification to existing information in the first database, and a deletion of existing information from the first database.

Claim 16. (original): The method of Claim 14, further comprising storing the first changes in at least one change table.

Claim 17. (original): The method of Claim 16, wherein storing the first changes in at least one change table comprises storing at least a portion of a row of information from the first database, a row identifier, and a time stamp in the change table.

Claim 18. (original): The method of Claim 17, wherein storing at least a portion of a row of information from the first database in the change table comprises storing an entire row of information in the change table, the first change affecting at least one entry in the row.

Claim 19. (original): The method of Claim 14, further comprising consolidating the first changes that affect information in one row of information from the first database.

Claim 20. (original): The method of Claim 19, wherein consolidating the first changes comprises: combining first changes that affect different entries in the row; and deleting at least one of a plurality of first changes that affect one or more common entries in the row.

Claim 21. (original): The method of Claim 14, further comprising: determining if any of the second changes conflict with the first changes; and deleting one of the first and second changes that conflict.

Claim 22. (original): The method of Claim 1, wherein the web servers are operable to execute stateless applications.

Claim 23. (original): The method of Claim 1, wherein the client is operable to use the first network address until the first network address fails and then use the second network address.

Claim 24. (original): A system for multi-site clustering in a network, comprising:
at least one computer readable medium; and
logic encoded on the computer readable medium and operable when executed to:

receive a domain name from a client;

identify a first network address and a second network address associated with the domain name, the first network address associated with a first site that includes a load balancer coupled to a plurality of web servers, the second network address associated with a second site that includes a domain name server;

determine whether the first site is available to serve the client;

determine whether the second site is available to serve the client;

communicate the first network address followed by the second network address to the client based at least partially on a determination that the first site and the second site are available; and

communicate the second network address to the client based at least partially on a determination that the first site is not available.

Claim 25. (original): The system of Claim 24, wherein the logic is operable to determine whether the first site is available by receiving an indication from the load balancer identifying whether at least one of the web servers is available.

Claim 26. (original): The system of Claim 25, wherein the load balancer is operable to determine whether at least one of the web servers is available by:

communicating a message to the web servers; and

determining whether the web servers respond to the message.

Claim 27. (original): The system of Claim 24, wherein:
the web servers are coupled to a database; and
the logic is operable to determine whether the first site is available by receiving an indication from the load balancer identifying whether the database is available.

Claim 28. (original): The system of Claim 27, wherein the load balancer is operable to determine whether the database is available by:

instructing at least one of the web servers to execute a web page;
receiving the executed web page; and
determining if the executed web page contains at least one expected keyword.

Claim 29. (original): The system of Claim 24, wherein the logic is operable to determine whether the second site is available by determining whether the domain name server is available.

Claim 30. (original): The system of Claim 29, wherein the logic is operable to determine whether the domain name server is available by determining if at least one message has been received from the domain name server during each of a plurality of time periods.

Claim 31. (original): The system of Claim 24, wherein the logic is further operable to communicate the first network address to the client when the first site is available and the second site is not available.

Claim 32. (original): The system of Claim 24, wherein the logic is operable to communicate the first network address followed by the second network address to the client based also on at least one of a load placed on each of the first and second sites, a response time of each of the first and second sites, and a locality of the client to each of the first and second sites.

Claim 33. (original): The system of Claim 32, wherein the logic is further operable to communicate the second network address followed by the first network address to the client based on the determination that the first site and the second site are available and at least one of the load placed on each of the first and second sites, the response time of each of the first and second sites, and the locality of the client to each of the first and second sites.

Claim 34. (original): The system of Claim 24, wherein the first site includes a first database and the second site includes a second database, the databases operable to synchronize information in the first and second databases.

Claim 35. (original): The system of Claim 34, wherein the first database is operable to synchronize the information by:

- making first changes to the information in the first database;
- storing the first changes in at least one change table;
- communicating the first changes in the change table to the second database;
- receiving second changes from the second database; and
- replicating the second changes to the information in the first database.

Claim 36. (original): The system of Claim 35, wherein the first database is further operable to:

- combine first changes that affect different entries in the row;
- delete at least one of a plurality of first changes that affect one or more common entries in the row;
- determine if any of the second changes conflict with the first changes; and delete one of the first and second changes that conflict.

Claim 37. (original): The system of Claim 24, wherein the web servers are operable to execute stateless applications.

Claim 38. (original): The system of Claim 24, wherein the client is operable to use the first network 20 address until the first network address fails and then use the second network address.

Claim 39. (original): A system for multi-site clustering in a network, comprising:
a memory operable to store a first network address and a second network address
associated with a domain name; and

at least one processor operable to:

receive the domain name from a client;

identify the first network address and the second network address associated with
the domain name, the first network address associated with a first site that includes a load
balancer coupled to a plurality of web servers, the second network address associated
with a second site that includes a domain name server;

determine whether the first site is available to serve the client;

determine whether the second site is available to serve the client;

communicate the first network address followed by the second network address to
the client based at least partially on a determination that the first site and the second site
are available; and

communicate the second network address to the client based at least partially on a
determination that the first site is not available.

Claim 40. (original): The system of Claim 39, wherein the processor is operable to
determine whether the first site is available by receiving an indication from the load balancer
identifying whether at least one of the web servers is available.

Claim 41. (original): The system of Claim 40, wherein the load balancer is operable
to determine whether at least one of the web servers is available by:

communicating a message to the web servers; and

determining whether the web servers respond to the message.

Claim 42. (original): The system of Claim 39, wherein:
the web servers are coupled to a database; and
the processor is operable to determine whether the first site is available by receiving an
indication from the load balancer identifying whether the database is available.

Claim 43. (original): The system of Claim 42, wherein the load balancer is operable to determine whether the database is available by:

instructing at least one of the web servers to execute a web page;
receiving the executed web page; and
determining if the executed web page contains at least one expected keyword.

Claim 44. (original): The system of Claim 39, wherein the processor is operable to determine whether the second site is available by determining whether the domain name server is available.

Claim 45. (original): The system of Claim 44, wherein the processor is operable to determine whether the domain name server is available by determining if at least one message has been received from the domain name server during each of a plurality of time periods.

Claim 46. (original): The system of Claim 39, wherein the processor is further operable to communicate the first network address to the client when the first site is available and the second site is not available.

Claim 47. (original): The system of Claim 39, wherein the processor is operable to communicate the first network address followed by the second network address to the client based also on at least one of a load placed on each of the first and second sites, a response time of each of the first and second sites, and a locality of the client to each of the first and second sites.

Claim 48. (original): The system of Claim 47, wherein the processor is further operable to communicate the second network address followed by the first network address to the client based on the determination that the first site and the second site are available and at least one of the load placed on each of the first and second sites, the response time of each of the first and second sites, and the locality of the client to each of the first and second sites.

Claim 49. (original): The system of Claim 39, wherein the first site includes a first database and the second site includes a second database, the databases operable to synchronize information in the first and second databases.

Claim 50. (original): The system of Claim 49, wherein the first database is operable to synchronize the information by:

- making first changes to the information in the first database;
- storing the first changes in at least one change table;
- communicating the first changes in the change table to the second database; receiving second changes from the second database; and
- replicating the second changes to the information in the first database.

Claim 51. (original): The system of Claim 50, wherein the first database is further operable to: combine first changes that affect different entries in the row;

- delete at least one of a plurality of first changes that affect one or more common entries in the row;

- determine if any of the second changes conflict with the first changes; and delete one of the first and second changes that conflict.

Claim 52. (original): The system of Claim 39, wherein the web servers are operable to execute stateless applications.

Claim 53. (original): The system of Claim 39, wherein the client is operable to use the first network address until the first network address fails and then use the second network address.

Claim 54. (original): A method for multi-site clustering in a network, comprising:
receiving a domain name from a client at a first domain name server;

identifying a first network address and a second network address associated with the domain name, the first network address associated with a first site that includes a load balancer coupled to a plurality of web servers and a first database, at least one of the web servers operable to execute stateless applications, the second network address associated with a second site that includes a second domain name server and a second database;

determining whether the first site is available to serve the client by:

communicating a message to the web servers;

determining whether the web servers respond to the message;

instructing at least one of the web servers to execute a web page;

receiving the executed web page; and

determining if the executed web page contains at least one expected keyword;

determining whether the second site is available to serve the client by determining if at least one message has been received at the first domain name server from the second domain name server during each of a plurality of time periods;

communicating the first network address followed by the second network address to the client based on a determination that the first site and the second site are available and on at least one of a load placed on each of the first and second sites, a response time of each of the first and second sites, and a locality of the client to each of the first and second sites; and

communicating the second network address to the client based at least partially on a determination that the first site is not available.

Claim 55. (original): A system for multi-site clustering in a network, comprising:
at least one computer readable medium; and
logic encoded on the computer readable medium and operable when executed to:

receive a domain name from a client;

identify a first network address and a second network address associated with the domain name, the first network address associated with a first site that includes a load balancer coupled to a plurality of web servers and a first database, at least one of the web servers operable to execute stateless applications, the second network address associated with a second site that includes a domain name server and a second database;

determine whether the first site is available to serve the client by:

communicating a message to the web servers;

determining whether the web servers respond to the message;

instructing at least one of the web servers to execute a web page; receiving the executed web page; and

determining if the executed web page contains at least one expected keyword;

determine whether the second site is available to serve the client by determining if at least one message has been received from the domain name server during each of a plurality of time periods;

communicate the first network address followed by the second network address to the client based on a determination that the first site and the second site are available and on at least one of a load placed on each of the first and second sites, a response time of each of the first and second sites, and a locality of the client to each of the first and second sites; and

communicate the second network address to the client based at least partially on a determination that the first site is not available.

Claim 56. (original): A system for multi-site clustering in a network, comprising:
a memory operable to store a first network address and a second network address associated with a domain name; and
at least one processor operable to:
 receive the domain name from a client;
 identify the first network address and the second network address associated with the domain name, the first network address associated with a first site that includes a load balancer coupled to a plurality of web servers and a first database, at least one of the web servers operable to execute stateless applications, the second network address associated with a second site that includes a domain name server and a second database;
 determine whether the first site is available to serve the client by:
 communicating a message to the web servers;
 determining whether the web servers respond to the message; instructing at least one of the web servers to execute a web page;
 receiving the executed web page; and
 determining if the executed web page contains at least one expected keyword;
 determine whether the second site is available to serve the client by determining if at least one message has been received from the domain name server during each of a plurality of time periods;
 communicate the first network address followed by the second network address to the client based on a determination that the first site and the second site are available and on at least one of a load placed on each of the first and second sites, a response time of each of the first and second sites, and a locality of the client to each of the first and second sites; and
 communicate the second network address to the client based at least partially on a determination that the first site is not available.